

Invertible Orientation Scores of 3D Images

Michiel Janssen¹, Remco Duits^{1,2}, and Marcel Breeuwer²

Eindhoven University of Technology, The Netherlands,

¹Department of Mathematics and Computer Science,

²Department of Biomedical Engineering.

M.H.J.Janssen@tue.nl, R.Duits@tue.nl, M.Breeuwer@tue.nl

Abstract. The enhancement and detection of elongated structures in noisy image data is relevant for many biomedical applications. To handle complex crossing structures in 2D images, 2D orientation scores $U : \mathbb{R}^2 \times S^1 \rightarrow \mathbb{R}$ were introduced, which already showed their use in a variety of applications. Here we extend this work to 3D orientation scores $U : \mathbb{R}^3 \times S^2 \rightarrow \mathbb{R}$. First, we construct the orientation score from a given dataset, which is achieved by an invertible coherent state type of transform. For this transformation we introduce 3D versions of the 2D cake-wavelets, which are complex wavelets that can simultaneously detect oriented structures and oriented edges. For efficient implementation of the different steps in the wavelet creation we use a spherical harmonic transform. Finally, we show some first results of practical applications of 3D orientation scores.

Key words: Orientation Scores, Reproducing Kernel Spaces, 3D Wavelet Design, Scale Spaces on SE(3), Coherence Enhancing Diffusion on SE(3)

1 Introduction

The enhancement and detection of elongated structures is important in many biomedical image analysis applications. These tasks become problematic when multiple elongated structures cross or touch each other in the data. In these cases it is useful to decompose an image in local orientations by constructing an orientation score. In the orientation score, we extend the domain of the data to include orientation in order to separate the crossing or touching structures (Fig. 1). From 3D data $f : \mathbb{R}^3 \rightarrow \mathbb{R}$ we construct a 3D orientation score $U : \mathbb{R}^3 \times S^2 \rightarrow \mathbb{R}$, in a similar way as is done for the more common case of 2D data $f : \mathbb{R}^2 \rightarrow \mathbb{R}$ and 2D orientation score $U : \mathbb{R}^2 \times S^1 \rightarrow \mathbb{R}$. Next, we consider operations on orientation scores, and process our data via orientation scores (Fig. 2). For such operations it is important that the orientation score transform is invertible, in a well-posed manner. In comparison to continuous wavelet transforms on the group of 3D rotations, translations and scalings, we use all scales simultaneously and exclude the scaling group from the wavelet transform and its adjoint, yielding a coherent state type of transform [1], see App.A. This makes it harder to design appropriate wavelets, but has the computational advantage of only needing a single scale transformation.

The 2D orientation scores have already showed their use in a variety of applications. In [11, 17] the orientation scores were used to perform crossing-preserving coherence-enhancing diffusions. These diffusions greatly reduce the noise in the data, while preserving the elongated crossing structures. Next to these generic enhancement techniques, the orientation scores also showed their use in retinal vessel segmentation [3], where they were used to better handle crossing vessels in the segmentation procedure.

To perform detection and enhancement operations on the orientation score, we first need to transform a given greyscale image or 3D dataset to an orientation score in an invertible way. In previous works various wavelets were introduced to perform a 2D orientation score transform. Some of these wavelets did not allow for an invertible transformation (e.g. Gabor wavelets [15]). A wavelet that allows an invertible transformation was proposed by Kalitzin [14]. A generalization of these wavelets was found by Duits [8] who derived a unitarity result and expressed the wavelets in a basis of eigenfunctions of the harmonic oscillator. This type of wavelet was also extended to 3D. This wavelet however has some unwanted properties such as poor spatial localization (oscillations) and the fact that the maximum of the wavelet did not lie at its center [8, Fig. 4.11]. In [8] a class of cake-wavelets were introduced, that have a cake-piece shaped form in the Fourier domain (Fig. 5). The cake-wavelets simultaneously detect oriented structures and oriented edges by constructing a complex orientation score $U : \mathbb{R}^2 \times S^1 \rightarrow \mathbb{C}$. Because the different cake-wavelets cover the full Fourier spectrum, invertibility is guaranteed.

In this paper we propose an extension of the 2D cake-wavelets to 3D. First, we discuss the theory of invertible orientation score transforms. Then we construct 3D cake-wavelets and give an efficient implementation using a spherical harmonic transform. Finally we mention two application areas for 3D orientation scores and show some preliminary results for both of them. In the first application, we present a practical proof of concept of a natural extension of the crossing preserving coherence enhancing diffusion on invertible orientation scores (CEDOS) [11] to the 3D setting. Compared to the original idea of coherence enhancing diffusion acting directly on image-data [18, 4, 5] we have the advantage of preserving crossings. Diffusions on $SE(3)$ have been studied in previous SSVM-articles, see e.g. [6], but the full generalization of CEDOS to 3D was never established.

2 Invertible Orientation Scores

An invertible orientation score $\mathcal{W}_\psi[f] : \mathbb{R}^3 \times S^2 \rightarrow \mathbb{C}$ is constructed from a given ball-limited 3D dataset $f \in \mathbb{L}_2^\varrho(\mathbb{R}^3) = \{f \in \mathbb{L}_2(\mathbb{R}^3) | \text{supp}(\mathcal{F}f) \subset B_{0,\varrho}\}$, with $\varrho > 0$ by correlation \star with an anisotropic kernel

$$(\mathcal{W}_\psi[f])(\mathbf{x}, \mathbf{n}) = (\overline{\psi_\mathbf{n}} \star f)(\mathbf{x}) = \int_{\mathbb{R}^3} \overline{\psi_\mathbf{n}(\mathbf{x}' - \mathbf{x})} f(\mathbf{x}') d\mathbf{x}', \quad (1)$$

where $\psi \in \mathbb{L}_2(\mathbb{R}^3) \cap \mathbb{L}_1(\mathbb{R}^3)$ is a wavelet aligned with and rotationally symmetric around the z -axis, and $\psi_\mathbf{n}(\mathbf{x}) = \psi(\mathbf{R}_\mathbf{n}^T \mathbf{x}) \in \mathbb{L}_2(\mathbb{R}^3)$ the rotated wavelet aligned

with \mathbf{n} . Here $\mathbf{R}_\mathbf{n}$ is any rotation which rotates the z -axis onto \mathbf{n} where the specific choice of rotation does not matter because of the rotational symmetry of ψ . The overline denotes a complex conjugate. The exact reconstruction formula for this transformation is

$$\begin{aligned} f(\mathbf{x}) &= (\mathcal{W}_\psi^{-1}[\mathcal{W}_\psi[f]])(\mathbf{x}) \\ &= \mathcal{F}_{\mathbb{R}^3}^{-1} \left[M_\psi^{-1} \mathcal{F}_{\mathbb{R}^3} \left[\tilde{\mathbf{x}} \mapsto \int_{S^2} (\check{\psi}_\mathbf{n} * \mathcal{W}_\psi[f](\cdot, \mathbf{n}))(\tilde{\mathbf{x}}) d\sigma(\mathbf{n}) \right] \right] (\mathbf{x}), \end{aligned} \quad (2)$$

with $\mathcal{F}_{\mathbb{R}^3}$ the Fourier transform on \mathbb{R}^3 given by $(\mathcal{F}f)(\boldsymbol{\omega}) = (2\pi)^{-\frac{3}{2}} \int_{\mathbb{R}^3} e^{-i\boldsymbol{\omega} \cdot \mathbf{x}} f(\mathbf{x}) d\mathbf{x}$ and $\check{\psi}_\mathbf{n}(\mathbf{x}) = \psi_\mathbf{n}(-\mathbf{x})$. In fact \mathcal{W}_ψ is a unitary mapping on to a reproducing kernel space, see App. A. The function $M_\psi : \mathbb{R}^3 \rightarrow \mathbb{R}^+$ is given by

$$M_\psi(\boldsymbol{\omega}) = (2\pi)^{\frac{3}{2}} \int_{S^2} |\mathcal{F}_{\mathbb{R}^3}[\psi_\mathbf{n}](\boldsymbol{\omega})|^2 d\sigma(\mathbf{n}). \quad (3)$$

The function M_ψ quantifies the stability of the inverse transformation [8], since $M_\psi(\boldsymbol{\omega})$ specifies how well frequency component $\boldsymbol{\omega}$ is preserved by the cascade of construction and reconstruction when M_ψ^{-1} would not be included in Eq. (2). An exact reconstruction is possible as long as

$$\exists_{M>0, \delta>0} \quad 0 < \delta \leq M_\psi(\boldsymbol{\omega}) \leq M < \infty, \quad \text{for all } \boldsymbol{\omega} = B_{0,\varrho}. \quad (4)$$

In practice it is best to aim for $M_\psi(\boldsymbol{\omega}) \approx 1$, in view of the condition number of $W_\psi : \mathbb{L}_2^\varrho(\mathbb{R}^3) \rightarrow \mathbb{L}_2^\varrho(\mathbb{R}^3 \times S^2)$ with $W_\psi f = \mathcal{W}_\psi f$. Also, when $M_\psi(\boldsymbol{\omega}) = 1$ we have \mathbb{L}_2 -norm preservation

$$\|f\|_{\mathbb{L}_2(\mathbb{R}^3)}^2 = \|\mathcal{W}_\psi f\|_{\mathbb{L}_2(\mathbb{R}^3 \times S^2)}^2, \quad \text{for all } f \in \mathbb{L}_2^\varrho(\mathbb{R}^3), \quad (5)$$

and Eq. (2) simplifies to $f(\mathbf{x}) = \int_{S^2} (\check{\psi}_\mathbf{n} * \mathcal{W}_\psi[f](\cdot, \mathbf{n}))(\mathbf{x}) d\sigma(\mathbf{n})$. We can further simplify the reconstruction for wavelets for which $(2\pi)^{\frac{3}{2}} \int_{S^2} \mathcal{F}_{\mathbb{R}^3}[\psi_\mathbf{n}](\boldsymbol{\omega}) d\sigma(\mathbf{n}) \approx 1$, where the reconstruction formula simplifies to an integration over orientations

$$f(\mathbf{x}) \approx \int_{S^2} \mathcal{W}_\psi f(\mathbf{x}, \mathbf{n}) d\sigma(\mathbf{n}). \quad (6)$$




Fig. 1. 2D Orientation score for an exemplary image. In the orientation score crossing structures are disentangled because the different structures have a different orientation.




Fig. 2. A schematic view of image processing via invertible orientation scores.




Fig. 3. Creating a 3D orientation score. Top: The data f is correlated with an oriented filter $\psi_{\mathbf{e}_x}$ to detect structures aligned with the filter orientation \mathbf{e}_x . Bottom left: This is repeated for a discrete set of filters with different orientations. Bottom right: The collection of 3D datasets constructed by correlation with the different filters is an orientation score and is visualized by placing a 3D dataset on a number of orientations.

2.1 Discrete Invertible Orientation Score Transformation

In the previous section, we considered a continuous orientation score transformation. In practice, we have only a finite number of orientations. To determine this discrete set of orientations we uniformly sample the sphere using platonic solids and/or refine this using tessellations of the platonic solids.

Assume we have a number N_o of orientations $\mathcal{V} = \{\mathbf{n}_1, \mathbf{n}_2, \dots, \mathbf{n}_{N_o}\} \subset S^2$, and define the discrete invertible orientation score $\mathcal{W}_{\psi}^d[f] : \mathbb{R}^3 \times \mathcal{V} \rightarrow \mathbb{C}$ by

$$(\mathcal{W}_{\psi}^d[f])(\mathbf{x}, \mathbf{n}_i) = (\overline{\psi_{\mathbf{n}_i}} \star f)(\mathbf{x}). \quad (7)$$

The exact reconstruction formula is in the discrete setting given by

$$\begin{aligned} f(\mathbf{x}) &= ((\mathcal{W}_{\psi}^d)^{-1}[\mathcal{W}_{\psi}^d[f]])(\mathbf{x}) \\ &= \mathcal{F}_{\mathbb{R}^3}^{-1} \left[(M_{\psi}^d)^{-1} \mathcal{F}_{\mathbb{R}^3} \left[\tilde{\mathbf{x}} \rightarrow \sum_{i=1}^{N_o} (\check{\psi}_{\mathbf{n}_i} \star \mathcal{W}_{\psi}^d[f](\cdot, \mathbf{n}_i))(\tilde{\mathbf{x}}) d\sigma(\mathbf{n}_i) \right] \right] (\mathbf{x}), \end{aligned} \quad (8)$$

with $d\sigma(\mathbf{n}_i)$ the discrete spherical area measure which for reasonably uniform spherical sampling can be approximated by $d\sigma(\mathbf{n}_i) \approx \frac{4\pi}{N_o}$, and

$$M_\psi^d(\boldsymbol{\omega}) = (2\pi)^{\frac{3}{2}} \sum_{i=1}^{N_o} |\mathcal{F}_{\mathbb{R}^3}[\psi_{\mathbf{n}_i}](\boldsymbol{\omega})|^2 d\sigma(\mathbf{n}_i). \quad (9)$$

Again, an exact reconstruction is possible iff $0 < \delta \leq M_\psi^d(\boldsymbol{\omega}) \leq M < \infty$.

3 3D Cake-Wavelets

A class of 2D cake-wavelets, see [8], was successfully used for the 2D orientation score transformation. We now generalize these 2D cake-wavelets to 3D cake-wavelets. Our 3D transformation using the 3D cake-wavelets should fulfill a set of requirements, compare [11] :

1. The orientation score should be constructed for a finite number (N_o) of orientations.
2. The transformation should be invertible and all frequencies should be transferred equally to the orientation score domain ($M_\psi^d \approx 1$).
3. The kernel should be strongly directional.
4. The kernel should be polar separable in the Fourier domain, i.e., $(\mathcal{F}\psi)(\boldsymbol{\omega}) = g(\rho)h(\theta, \phi)$, with $\boldsymbol{\omega} = (\omega_x, \omega_y, \omega_z) = (\rho \sin \theta \cos \phi, \rho \sin \theta \sin \phi, \rho \cos \theta)$. Because by definition the wavelet ψ has rotational symmetry around the z -axis we have $h(\theta, \phi) = h(\theta)$.
5. The kernel should be localized in the spatial domain, since we want to pick up local oriented structures.
6. The real part of the kernel should detect oriented structures and the imaginary part should detect oriented edges. The constructed oriented score is therefore a complex orientation score.

3.1 Construction of Line and Edge Detectors

We now discuss the procedure used to make 3D cake-wavelets. According to requirement 4 we only consider polar separable wavelets in the Fourier domain, so that $(\mathcal{F}\psi)(\boldsymbol{\omega}) = g(\rho)h(\theta)$. For the radial function $g(\rho)$ we use, as in [11],

$$g(\rho) = \mathcal{M}_N(\rho^2 t^{-1}) = e^{-\frac{\rho^2}{t}} \sum_{k=0}^N \frac{(\rho^2 t^{-1})^k}{k!}, \quad (10)$$

which is a Gaussian function with scale parameter t multiplied by the Taylor approximation of its reciprocal to order N to ensure a slower decay. This function should go to 0 when ρ tends to the Nyquist frequency ρ_N . Therefore the inflection point of this function is fixed at $\gamma \rho_N$ with $0 \ll \gamma < 1$ by setting $t = \frac{2(\gamma \rho_N)^2}{1+2N}$. In practice we have $\varrho = \rho_N$, and because radial function g causes

M_ψ^d to become really small when coming close to the Nyquist frequency, reconstruction Eq.(8) becomes unstable. We solve this by either using approximate reconstruction Eq.(6) or by replacing $M_\psi^d \rightarrow \max(M_\psi^d, \epsilon)$, with ϵ small. Both make the reconstruction stable at the cost of not completely reconstructing the highest frequencies which causes some additional blurring.

We now need to find an appropriate angular part \mathfrak{h} for the cake-wavelets. First, we specify an orientation distribution $A : S^2 \rightarrow \mathbb{R}^+$, which determines what orientations the wavelet should measure. To satisfy requirement 3 this function should be a localized spherical window, for which we propose a B-spline $A(\theta, \phi) = B^k\left(\frac{\theta}{s_\theta}\right)$, with $s_\theta > 0$ and B^k the k th order B-spline given by

$$B^k(x) = (B^{k-1} * B^0)(x), \quad B^0(x) = \begin{cases} 1 & \text{if } -\frac{1}{2} < x < \frac{1}{2} \\ 0 & \text{otherwise} \end{cases}. \quad (11)$$

The parameter s_θ determines the trade-off between requirements 2 and 3, where higher values give a more uniform M_ψ^d at the cost of less directionality.

First consider setting $h = A$ so that ψ has compact support within a convex cone in the Fourier domain. The real part of the corresponding wavelet would however be a plate detector and not a line detector (Fig. 4). The imaginary part is already an oriented edge detector, and so we set

$$\mathfrak{h}_{Im}(\phi) = A(\theta, \phi) - A(\pi - \theta, \phi + \pi) = B^k\left(\frac{\theta}{s_\theta}\right) - B^k\left(\frac{\pi - \theta}{s_\theta}\right), \quad (12)$$

where the real part of the earlier found wavelet vanishes by anti-symmetrization of the orientation distribution A while the imaginary part remains. As to the construction of h_{Re} , there is the general observation that we detect a structure that is perpendicular to the shape in the Fourier domain, so for line detection we should aim for a plane detector in the Fourier domain. To achieve this we apply the Funk transform to A , and we define

$$h_{Re}(\theta, \phi) = FA(\theta, \phi) = \int_{S_p(\mathbf{n}(\theta, \phi))} A(\mathbf{n}') ds(\mathbf{n}'), \quad (13)$$

where integration is performed over $S_p(\mathbf{n})$ denoting the great circle perpendicular to $\mathbf{n}(\theta, \phi) = (\sin \theta \cos \phi, \sin \theta \sin \phi, \cos \theta)$. This transformation preserves the symmetry of A , so we have $h_{Re}(\theta, \phi) = \mathfrak{h}_{Re}(\theta)$. Thus, we finally set

$$\mathfrak{h}(\theta) = \mathfrak{h}_{Re}(\theta) + \mathfrak{h}_{Im}(\theta). \quad (14)$$

For an overview of the transformations see Fig. 5.

3.2 Efficient Implementations Via Spherical Harmonics

In Subsection 3.1 we defined the real part and the imaginary part of the wavelets in terms of a given orientation distribution. In order to efficiently implement the

various transformations (e.g. Funk transform), and to create the various rotated versions of the wavelet we express our orientation distribution A in a spherical harmonic basis $\{Y_l^m\}$ up to order L :

$$A(\theta, \phi) = \sum_{l=0}^L \sum_{m=-l}^l c_{l,m} Y_l^m(\theta, \phi), \quad L \in \mathbb{N}. \quad (15)$$

Because of the rotational symmetry around the z -axis, we only need the spherical harmonics with $m = 0$, i.e., $A(\theta, \phi) = \sum_{l=0}^L c_{l,0} Y_l^0(\theta, \phi)$. For determining the spherical harmonic coefficients we use the pseudo-inverse of the discretized inverse spherical harmonic transform (see [9, Section 7.1]), with discrete orientations given by an icosahedron of tesselation order 15.

Funk Transform According to [7], the Funk transform of a spherical harmonic equals

$$FY_l^m(\theta, \phi) = \int_{S_p(\mathbf{n}(\theta, \phi))} Y_l^m(\mathbf{n}') ds(\mathbf{n}') = 2\pi P_l(0) Y_l^m(\theta, \phi), \quad (16)$$

with $P_l(0)$ the Legendre polynomial of degree l evaluated at 0. We can therefore apply the Funk transform to a function expressed in a spherical harmonic basis by a simple transformation of the coefficients $c_l^m \rightarrow 2\pi P_l(0) c_l^m$.

Anti-Symmetrization We have $Y_l^m(\pi - \theta, \phi + \pi) = (-1)^l Y_l^m(\theta, \phi)$. We therefore anti-symmetrize the orientation distribution Eq.(12) via $c_l^m \rightarrow (1 - (-1)^l) c_l^m$.

Making Rotated Wavelets To make the rotated versions $\psi_{\mathbf{n}}$ of wavelet ψ we have to find $h_{\mathbf{n}}$ in $\Psi_{\mathbf{n}} = g(\rho) h_{\mathbf{n}}(\theta, \phi)$. To achieve this we use the steerability of the spherical harmonic basis. Spherical harmonics rotate according to the irreducible representations of the $SO(3)$ group $D_{m,m'}^l(\alpha, \beta, \gamma)$ (Wigner-D functions)

$$\mathcal{R}_{\mathbf{R}_{\alpha,\beta,\gamma}} Y_l^m(\theta, \phi) = \sum_{m'=l}^l D_{m,m'}^l(\alpha, \beta, \gamma) Y_l^{m'}(\theta, \phi). \quad (17)$$

Here α, β and γ denote the Euler angles with counterclockwise rotations, i.e., $\mathbf{R} = \mathbf{R}_{\mathbf{e}_z, \alpha} \mathbf{R}_{\mathbf{e}_y, \beta} \mathbf{R}_{\mathbf{e}_z, \gamma}$. This gives

$$h_{\mathbf{n}}(\theta, \phi) = \mathcal{R}_{\mathbf{R}_{\alpha,\beta,\gamma}} h(\theta, \phi) = \sum_{l=0}^L \sum_{m=-l}^l \sum_{m'=-l}^l a_{l,m} D_{m,m'}^l(\alpha, \beta, \gamma) Y_l^{m'}(\theta, \phi). \quad (18)$$

Because both anti-symmetrization and Funk transform preserve the rotational symmetry of A , we have $h(\theta, \phi) = \sum_{l=0}^L a_{l,0} Y_l^0(\theta, \phi)$, and Eq. (18) reduces to

$$h_{\mathbf{n}}(\theta, \phi) = \sum_{l=0}^L \sum_{m'=-l}^l a_{l,0} D_{0,m'}^l(0, \beta, \gamma) Y_l^{m'}(\theta, \phi). \quad (19)$$




Fig. 4. When directly setting orientation distribution A as angular part of the wavelet h we construct plate detectors. From left to right: Orientation distribution A , wavelet in the Fourier domain, the plate detector (real part) and the edge detector (imaginary part). Orange: Positive iso-contour. Blue: Negative iso-contour. Parameters used: $L = 16$, $s_\theta = 0.6$, $k = 2$, $N = 20$, $\gamma = 0.8$ and evaluated on a grid of $51 \times 51 \times 51$ pixels.




Fig. 5. Cake-Wavelets. Top: 2D cake-wavelets. From left to right: Illustration of the Fourier domain coverage, the wavelet in the Fourier domain and the real and imaginary part of the wavelet in the spatial domain. [3]. Bottom: 3D cake-wavelets. Overview of the transformations used to construct the wavelets from a given orientation distribution. Upper part: The wavelet according to Eq. (12). Lower part: The wavelet according to Eq. (13). IFT: Inverse Fourier Transform. Parameters used: $L = 16$, $s_\theta = 1.05$, $k = 2$, $N = 20$, $\gamma = 0.8$ and evaluated on a grid of $31 \times 31 \times 31$ pixels.

4 Applications

4.1 Adaptive Crossing Preserving Flows

We now use the invertible orientation score transformation to perform data-enhancement according to Fig. 2. Because $\mathbb{R}^3 \times S^2$ is not a Lie group, it is common practice to embed the space of positions and orientations in the Lie group of positions and rotations SE(3) by setting

$$\tilde{U}(\mathbf{x}, \mathbf{R}) = U(\mathbf{x}, \mathbf{R} \cdot \mathbf{e}_z), \quad U(\mathbf{x}, \mathbf{n}) = \tilde{U}(\mathbf{x}, \mathbf{R}_n), \quad (20)$$

with \mathbf{R}_n any rotation for which $\mathbf{R}_n \cdot \mathbf{e}_z = \mathbf{n}$. This holds in particular for orientation scores $U = \mathcal{W}_\psi f$. The operations Φ which we consider are scale spaces on SE(3) (diffusions), and are given by $\Phi = \Phi_t$ with

$$\Phi_t(U)(\mathbf{y}, \mathbf{n}) = \tilde{W}(y, \mathbf{R}_n, t). \quad (21)$$

Here \tilde{W} is the solution of

$$\frac{\partial \tilde{W}}{\partial t}(g, t) = \sum_{i,j=1}^6 \mathcal{A}_i|_g D_{ij} \mathcal{A}_j|_g \tilde{W}(g, t), \quad \tilde{W}|_{t=0} = \widetilde{\mathcal{W}_\psi[f]}, \quad (22)$$

where in coherence enhancing diffusion on orientation scores (CEDOS) D_{ij} is adapted locally to data $\widetilde{\mathcal{W}_\psi[f]}$ based on exponential curve fits (see [10]), and with $\mathcal{A}_i|_{g=(\mathbf{x}, \mathbf{R})} = (L_g)_* \mathcal{A}_i|_e$ the left-invariant vector fields on SE(3), for motivation and details see [9]. Furthermore D_{ij} is chosen such that equivalence relation Eq. (20) is maintained for \tilde{W} . These operations are already used without adaptivity in the field of diffusion weighted MRI, where similar data (of the type $\mathbb{R}^3 \times S^2 \rightarrow \mathbb{R}^+$) is enhanced [9]. We then obtain Euclidean invariant image processing via

$$\Upsilon f = \mathcal{W}_\psi^{*,ext} \circ \Phi \circ \mathcal{W}_\psi f = \mathcal{W}_\psi^* \circ \mathbb{P}_\psi \Phi \circ \mathcal{W}_\psi f \quad (23)$$

which includes inherent projection \mathbb{P}_ψ of orientation scores, even if $\Phi = \Phi_t$ maps outside of the space of orientation scores in the embedding space (see App. A). Below we show some preliminary results of these flows that enhance the elongated structures while preserving the crossing, Fig. 6 and Fig. 7.

4.2 3D Vessel Tracking in Magnetic Resonance Angiography (MRA) Data

We use the 3D orientation scores to extend the earlier work on 2D vessel segmentation via invertible orientation scores [3] to 3D vessel segmentation in MRA-data. Even though true crossing structures hardly appear in 3D data, we do encounter vessels touching other vessels/structures. The orientation scores also allow us to better handle complex structures, such as bifurcations. In Fig. 8 we show some first results of the vessel segmentation algorithm.




Fig. 6. Adaptive Crossing Preserving Flows. From left to right 3D visualization of artificial data, slice of data, slice of (data + Gaussian noise), slice of enhanced data. For the orientation score transformation we use: $N_0 = 42, s_\theta = 0.7, k = 2, N = 20, \gamma = 0.85, L = 16$ evaluated on a grid of $21 \times 21 \times 21$ pixels. We use approximate reconstruction Eq.(8), and for diffusion we set $t = 10$. For the choice of D_{ij} in CEDOS, see [10].




Fig. 7. Adaptive Crossing Preserving Flows combined with soft thresholding $\Phi(U)(\mathbf{x}, \mathbf{n}) = |U(\mathbf{x}, \mathbf{n})|^{1.5} \operatorname{sgn}(U(\mathbf{x}, \mathbf{n}))$ on data containing the Adam Kiewitzc vessel. From left to right: Slice of data, data after soft thresholding, data after CEDOS, data after CEDOS followed by soft thresholding. For parameters see Fig.6, but now $t = 5$.




Fig. 8. MRA vessel segmentation via invertible orientation scores.

5 Conclusion

We have extended 2D cake-wavelets to 3D cake-wavelets, which can be used for a 3D invertible orientation score transformation. Efficient implementation for calculating the wavelets via spherical harmonics were introduced. The developed transformation allows us to consider all kinds of enhancement operations via orientation scores such as the adaptive crossing preserving flows which we are currently working on. Next to data-enhancement we also showed some first results of 3D vessel segmentation using 3D orientation scores.

Acknowledgements. We thank Dr. A.J.E.M. Janssen for advice on the presentation of this paper. The research leading to these results has received funding from the European Research Council under the European Community's Seventh Framework Programme (FP7/2007-2013) / ERC grant *Lie Analysis*, agr. nr. 335555.

A Invertible Orientation Scores of 3D-images and Continuous Wavelet Theory

The continuous wavelet transform constructed by unitary irreducible representations of locally compact groups was first formulated by Grossman et al. [13]. Given a Hilbert space H and a unitary irreducible representation $g \mapsto \mathcal{U}_g$ of any locally compact group G in H , a non-zero vector $\psi \in H$ is called admissible if

$$C_\psi := \int_G \frac{|(\mathcal{U}_g\psi, \psi)|^2}{(\psi, \psi)_H} d\mu_G(g) < \infty, \quad (24)$$

where μ_G denotes the left-invariant Haar measure. Given an admissible vector ψ and a unitary representation of a locally compact group G in H , the Coherent State (CS) transform $W_\psi : H \rightarrow \mathbb{L}_2(G)$ is given by $(W_\psi[f])(g) = (\mathcal{U}_g\psi, f)_H$. W_ψ is an isometric transform onto a unique closed reproducing kernel space $\mathbb{C}_{K_\psi}^G$ with $K_\psi(g, g') = \frac{1}{C_\psi}(\mathcal{U}_g\psi, \mathcal{U}_{g'}\psi)_H$ as an \mathbb{L}_2 -subspace [1].

We distinguish between the isometric wavelet transform $W_\psi : \mathbb{L}_2^\rho(\mathbb{R}^3) \rightarrow \mathbb{L}_2(G)$ and the unitary wavelet transform $\mathcal{W}_\psi : \mathbb{L}_2^\rho(\mathbb{R}^3) \rightarrow \mathbb{C}_{K_\psi}^G$. We drop the formal requirement of \mathcal{U} being square-integrable and ψ being admissible in the sense of (24), and replace the requirement by (4), as it is not strictly needed in many cases. This includes our case of interest $G = SE(3)$ and its left-regular action on $\mathbb{L}_2(\mathbb{R}^3)$ where \mathcal{W}_ψ gives rise to an orientation score $\mathcal{W}_\psi f : \mathbb{R}^3 \times S^2 \rightarrow \mathbb{C}$

$$\mathcal{W}_\psi f(\mathbf{x}, \mathbf{n}) = \widetilde{\mathcal{W}_\psi f}(\mathbf{x}, \mathbf{R}_\mathbf{n}), \quad (25)$$

with $\mathbf{R}_\mathbf{n}$ any rotation mapping \mathbf{e}_z onto \mathbf{n} and ψ symmetric around the z -axis. Here the domain is the coupled space of positions and orientations: $\mathbb{R}^3 \times S^2 := SE(3)/(\{\mathbf{0}\} \times SO(2))$, cf. [9].

From the general theory of reproducing kernel spaces [8, Thm 18],[2] (where one does not even rely on the group structure), it follows that $\mathcal{W}_\psi : \mathbb{L}_2^\rho(\mathbb{R}^3) \rightarrow \mathbb{C}_{K_\psi}^{\mathbb{R}^3 \times S^2}$ is unitary, where $\mathbb{C}_{K_\psi}^{\mathbb{R}^3 \times S^2}$ denotes the abstract complex reproducing kernel space consisting of functions on $\mathbb{R}^3 \times S^2$ with reproducing kernel

$$K_{(\mathbf{y}, \mathbf{n})}(\mathbf{y}', \mathbf{n}') = (\mathcal{U}_{(\mathbf{y}, \mathbf{R}_\mathbf{n})}\psi, \mathcal{U}_{(\mathbf{y}', \mathbf{R}_\mathbf{n}')}\psi)_{\mathbb{L}_2(\mathbb{R}^3)}, \quad (26)$$

with left-regular representation $(\mathbf{y}, \mathbf{R}) \mapsto \mathcal{U}_{(\mathbf{y}, \mathbf{R})}\psi$ given by $(\mathcal{U}_{(\mathbf{y}, \mathbf{R})}\psi)(\mathbf{x}) = \psi(\mathbf{R}^T(\mathbf{x} - \mathbf{y}))$. Now, as the characterization of the inner product on $\mathbb{C}_{K_\psi}^{\mathbb{R}^3 \times S^2}$ is awkward [16], we provide a basic characterization next via the so-called M_ψ inner product. This is in line with the admissibility conditions in [12].

Theorem 1. *Let ψ be such that (4) holds. Then $\mathcal{W}_\psi : \mathbb{L}_2^\rho(\mathbb{R}^3) \rightarrow \mathbb{C}_{K_\psi}^{\mathbb{R}^3 \times S^2}$ is unitary, and we have*

$$(f, g)_{\mathbb{L}_2(\mathbb{R}^3)} = (\mathcal{W}_\psi f, \mathcal{W}_\psi g)_{M_\psi}, \quad (27)$$

where $(\mathcal{W}_\psi f, \mathcal{W}_\psi g)_{M_\psi} = (\mathcal{T}_{M_\psi}[\mathcal{W}_\psi f], \mathcal{T}_{M_\psi}[\mathcal{W}_\psi g])_{\mathbb{L}_2(\mathbb{R}^3 \times S^2)}$, with $[\mathcal{T}_{M_\psi}[U]](\mathbf{y}, \mathbf{n}) := \mathcal{F}^{-1} \left[\boldsymbol{\omega} \mapsto (2\pi)^{-3/4} M_\psi^{-1/2}(\boldsymbol{\omega}) \mathcal{F}[U(\cdot, \mathbf{n})](\boldsymbol{\omega}) \right](\mathbf{y})$.

Proof. We rely on [17, Thm 1], where we set $H = \mathbb{L}_2(\mathbb{R}^3)$. The rest follows by well posed restriction to the quotient $\mathbb{R}^3 \rtimes S^2$.

Corollary 1. *Let $M_\psi > 0$ on \mathbb{R}^3 . The space $\mathbb{C}_K^{\mathbb{R}^3 \times S^2}$ is a closed subspace of Hilbert space $\mathbb{H}_\psi \otimes \mathbb{L}_2(S^2)$, where $\mathbb{H}_\psi = \{f \in \mathbb{L}_2(\mathbb{R}^3) \mid M_\psi^{-\frac{1}{2}}\mathcal{F}[f] \in \mathbb{L}_2(\mathbb{R}^3)\}$, and projection of embedding space onto the space of orientation scores is given by $(\mathbb{P}_\psi(U))(\mathbf{y}, \mathbf{n}) = (K_{(\mathbf{n}, \mathbf{y})}, U)_{M_\psi} = (\mathcal{W}_\psi \mathcal{W}_\psi^{*,ext}(U))(\mathbf{y}, \mathbf{n})$, where $\mathcal{W}_\psi^{*,ext}$ is the natural extension of the adjoint to the embedding space.*

References

1. S.T. Ali. A general theorem on square-integrability: Vector coherent states. *J. Math. Phys.*, 39(8):3954, 1998.
2. S.T. Ali, J.-P. Antoine, and J.-P. Gazeau. *Coherent states, wavelets, and their generalizations*. Springer, 2014.
3. E. Bekkers and R. Duits. A multi-orientation analysis approach to retinal vessel tracking. *JMIV*, 2014.
4. B. Burgeth, S. Didas, and J. Weickert. A general structure tensor concept and coherence-enhancing diffusion filtering for matrix fields. In *Visualization and processing of tensor fields*, pages 305–324. Springer, Berlin, 2009.
5. B. Burgeth, L. Pizarro, S. Didas, and J. Weickert. 3D-Coherence-enhancing diffusion filtering for matrix fields. In *Mathematical methods for signal and image analysis and representation*, pages 49–63. Springer London, 2012.
6. E.J. Creusen, R. Duits, and T.C.J. Dela Haije. Numerical schemes for linear and non-linear enhancement of DW-MRI. *SSVM*, 1:14–25, 2012.
7. M. Descoteaux, E. Angelino, S. Fitzgibbons, and R. Deriche. Regularized, fast, and robust analytical Q-ball imaging. *MRM*, 58(3):497–510, September 2007.
8. R. Duits. *Perceptual organization in image analysis*. PhD thesis, Technische Universiteit Eindhoven, 2005.
9. R. Duits and E.M. Franken. Left-invariant diffusions on the space of positions and orientations and their application to crossing-preserving smoothing of HARDI images. *IJCV*, 92(3):231–264, March 2010.
10. R. Duits, M.H.J. Janssen, J. Hannink, and G.R. Sanguinetti. Locally Adaptive Frames in the Roto-Translation Group and their Applications in Medical Imaging. *arXiv preprint: arXiv:1502.08002*.
11. E.M. Franken and R. Duits. Crossing-preserving coherence-enhancing diffusion on invertible orientation scores. *IJCV*, 85(3):253–278, February 2009.
12. F. Führ. Abstract harmonic analysis of continuous wavelet transforms. *Lecture Notes in Mathematics*, 1863, 2005.
13. A. Grossmann, J. Morlet, and T. Paul. Transforms associated to square integrable group representations. I. General results. *J. Math. Phys.*, 26(10):2473, 1985.
14. S.N. Kalitzin, B.M. ter Haar Romeny, and M.A. Viergever. Invertible apertured orientation filters in image analysis. *IJCV*, 31:145–158, 1999.
15. Tai Sing Lee. Image representation using 2D Gabor wavelets. *IEEE TPAMI*, 18(10), 1996.
16. F.J.L. Martens. *Spaces of analytical functions on inductive/projective limits of Hilbert spaces*. PhD thesis, Technische Universiteit Eindhoven, 1988.
17. U. Sharma and R. Duits. Left-invariant evolutions of wavelet transforms on the similitude group. *accepted for publication ACHA, doi:10.1016/j.acha.2014.09.001*.
18. J. Weickert. Coherence-enhancing diffusion filtering. *IJCV*, 31:111–127, 1999.